Attorney Docket No. IBM/177
Confirmation No. 3237

PATENT

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE BOARD OF PATENT APPEALS
AND INTERFERENCES

Ex parte Robert Miller, Vicki Lynn Morey and Laurie Ann Williams

---

Appeal No. _____
Application No. 09/845,596

AMENDED APPEAL BRIEF

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant:	Robert Miller et al.	Art Unit:	2157
Application No.:	09/845,596	Examiner:	Sargon N. Nano
Filed:	April 30, 2001	Atty. Docket No.:	IBM/177

For: GROUP ACCESS PRIVATIZATION IN CLUSTERED COMPUTER SYSTEM

Mail Stop Appeal Brief - Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

AMENDED APPEAL BRIEF

I. REAL PARTY IN INTEREST

This application is assigned to International Business Machines Corporation, of Armonk, New York.

II. RELATED APPEALS AND INTERFERENCES

There are no related appeals or interferences.

III. STATUS OF CLAIMS

Claims 1-27 are pending in the Application, stand rejected, and are now on appeal.

IV. STATUS OF AMENDMENTS

There have been no amendments filed subsequent to the final rejection mailed November 4, 2005.

V. SUMMARY OF CLAIMED SUBJECT MATTER

Applicant's invention is generally directed to an apparatus, program product and method that rely on "cluster-private" group names to access groups that are resident in a clustered computer system. (Application, p. 4, lines 2-5). A clustered computer system, within the context of the invention, is a computer system where multiple computers, or nodes, are networked

together to present a single system image (Application, p. 1, lines 9-14). Clusters perform tasks through the performance of jobs running on each node, which may be logically organized together into a "group" to perform collective tasks, where each affiliated job is referred to as a "member" of the group (Application, p. 1, line 27 to p. 2, line 4).

As described at page 2, lines 12-20 of the Application, external access to a group, e.g., to enable external entities to send work requests to a group, have conventionally been supported by assigning a specific network address (e.g., a TCP/IP address) to the group, such that an external entity wishing to access a group can send a request to that specific address. This technique is sometimes called N+1 addressing, where N addresses are assigned to the N nodes in a group, plus one additional address for the group itself.

In conventional clustered computer systems, a name service is typically provided to map network addresses of groups to "group names". (Application, p. 2, lines 21-23). In many conventional designs, these mappings are stored in a network name server such as a directory name service (DNS) server, so that, should an entity desire to access another entity on a network, the accessing entity can resolve the name of the entity using the network name server, and then send a message to the network address returned by the server. Thus, in the case of an external access to a group, an entity wishing to send a request to the group typically resolves the group name through the network name server, and sends a message to the group address that is returned by the server. (Application, p. 2, line 29 to p. 3, line 5).

Applicants have found, however, that the use of an external name server in connection with accessing a group presents a number of problems. First, a significant concern is presented that a node or other entity outside of a cluster could send messages to a group that could interfere with the group's operation, which poses a significant security risk to the system. Second, in many instances, it may be desirable to implement multiple clusters, or cluster "instances", on a given clustered computer system or network, e.g., in a logically partitioned system where multiple cluster instances may execute concurrently in different logical computer systems that execute on the same physical system. Where multiple clusters exist, however, a limitation is presented in that the same group name cannot exist in each cluster, since a common name server

that cannot resolve a group name to different network addresses is often used. (Application, p. 3, lines 6-20).

Applicants have addressed this problem through the use of group names that are "cluster-private" in nature. Put another way, for a cluster-accessible group, all nodes capable of participating in a cluster are configured to map to the same cluster-private group name for that group, so that any authorized external user that has access to the cluster can access the group name and utilize the group name to initiate operations by the group, but any unauthorized user (e.g., a user on an unauthorized node), is restricted from accessing the group name or issuing requests to a cluster using that group name, and is thus restricted from accessing the group. (Application, p. 4, lines 2-11).

One manner of making a group name "cluster-private" is through the use of a group address data structure that is accessible only on nodes that participate in a cluster. In addition, it may be desirable to require group names to be resolved to group member addresses locally on each node that participates in a cluster. By doing so, the mapping of cluster-private group names to the addresses of the group members (e.g., the network addresses of the nodes upon which the group members are resident) can often be effectively "hidden" from nodes or users that are outside of the cluster. (Application, p. 4, lines 12-19).

The concept of a "cluster-private" group name is discussed in further detail at p. 7, lines 6-13, and is reproduced below for the convenience of the Board:

> A group name is cluster-private if the group name cannot be accessed outside of a node that participates in a particular cluster, or cluster "instance". A group name may be cluster-private, for example, if the group name is resident upon a node that participates in a cluster instance, and is accessible by jobs executing on that node, i.e., if the group name is local to the node. A group name may also be cluster-private, for example, if some form of authentication mechanism is utilized to restrict access to the group name only to nodes and/or jobs that can establish participation in a cluster.

By utilizing cluster-private group names, external access to a group by unauthorized entities is substantially restricted due to the inability of such entities to obtain the appropriate group name and/or make a request that properly identifies the group name, as well as the local

resolution of name-address mappings within each node, thus enhancing overall cluster security. (Application, p. 8, lines 17-22. In addition, it is possible to support the use of the same group name by multiple cluster instances on the same network without conflict, e.g., in a logically-partitioned system where multiple cluster instances execute on the same network. (Application, p. 8, lines 23-26).

One manner of implementing cluster-private group names, for example, utilizes the clustering infrastructure resident on each node that participates in a particular cluster to maintain a local mapping of group names to the addresses of the members of the group and/or the nodes upon which such members reside. As such, how messages are distributed to group members, and in particular, the addresses of those members, are effectively hidden from jobs or applications that access the clustering infrastructure. Therefore, rather than initiating an operation in a group by sending a request to a particular address, the request is sent to the clustering infrastructure and managed by the clustering infrastructure at a higher software layer and protocol than the low level network addresses that have conventionally been used to access cluster groups. Moreover, in many instances, each node, and typically the clustering infrastructure within each node, is capable of locally resolving the group name - group member address mapping. Thus, an external name server, as is otherwise required by conventional implementations, is often not required. (Application, p. 8, lines 3-16).

As such, in many instances, the network topology utilized by a clustered computer system is irrelevant for the purposes of supporting group operations, since distribution of group messages in response to an authorized external request is managed within the clustering infrastructure of each participating node. (Application, p. 8, lines 26-30).

Among the claims on appeal, claims 1, 15, 25 and 26 are independent. Support for the claimed subject matter in these claims may be found, for example, in the Application, at p. 4, lines 2-32, p. 7, lines 2-20, Fig. 3 and the accompanying description at p. 13, line 20 to p. 14, line 21, and Fig. 4 and the accompanying description at p. 14, line 22 to p. 15, line 32.

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

Claims 1-27 stand rejected under 35 U.S.C. § 102(e) as being anticipated by U.S. Patent No. 6,470,389 to Chung et al. (*Chung*).

VII. ARGUMENT

Applicants respectfully submit that the Examiner's rejections of claims 1-27 are not supported on the record, and should be reversed. Anticipation of a claim under 35 U.S.C. §102 requires that "each and every element as set forth in the claim is found, either expressly or inherently described, in a single prior art reference." Verdegaal Bros., Inc. v. Union Oil Co., 2 USPQ2d 1051, 1053 (Fed. Cir. 1987), *quoted in* In re Robertson, 49 USPQ2d 1949, 1950 (Fed. Cir. 1999). Absent express description, anticipation under inherency requires extrinsic evidence that makes it clear that "the missing descriptive matter is necessarily present in the thing described in the reference, and that it would be so recognized by persons of ordinary skill." Continental Can Co. v. Monsanto Co., 20 USPQ2d 1746, 1749 (Fed. Cir. 1991), *quoted in* In re Robertson at 1951. "Inherency, however, may not be established by probabilities or possibilities. The mere fact that a certain thing may result from a given set of circumstances is not sufficient." Continental Can at 1749, *quoted in* In re Robertson at 1951.

Applicants respectfully submit that *Chung* does not disclose the various features recited in claims 1-27, and as such, the rejections thereof should be reversed. Applicants will hereinafter address the various claims that are the subject of the Examiner's rejection in order, starting with the independent claims, and then addressing additional dependent claims reciting additional subject matter that is distinguishable from *Chung*. In some cases, specific discussions of particular claims are not made in the interests of streamlining the appeal. The omission of a discussion with respect to any particular claim, however, should not be interpreted as an acquiescence as to the merits of the Examiner's rejection of the claim, particularly with respect to claims reciting features that are addressed in connection with the rejections applied to other claims pending in the appeal.

Independent Claim 1

Claim 1 recites a method of accessing a group in a clustered computer system, where the clustered computer system includes a plurality of nodes, and where the group includes a plurality of members resident respectively on the plurality of nodes. The recited method includes the step of receiving an access request on a first node in the plurality of nodes, where the access request identifies a cluster-private group name associated with the group. The recited method also includes the step of processing the access request on the first node to initiate a group operation on at least a subset of the plurality of nodes that map to the cluster-private group name.

As such, claim 1 is directed, in part, to the concept of processing an access request in a clustered computer system where the request identifies a cluster-private group name associated with a group that is resident on the clustered computer system.

In rejecting claim 1, the Examiner relies on *Chung*, and specifically col. 7, lines 13-52 thereof. However, this cited passage merely discloses the concept of using a "ghost" IP address, which, as described at col. 7, lines 27-28, is "publicized" to a DNS system to enable a client to obtain the IP address based upon a domain name associated with a web site or other network service hosted by the cluster.

Interestingly, it is this precise configuration that Applicants describe in the Background section of the present Application, at page 2, line 15 to page 3, line 5:

> Conventionally, external access to a group has been supported through assigning a specific network address (e.g., a TCP/IP address) to the group, such that an external entity wishing to access a group can send a request to that specific address. This technique is sometimes called N+1 addressing, where N addresses are assigned to the N nodes in a group, plus one additional address for the group itself. . . As with other conventional network addressing protocols, typically a name service is provided in a conventional clustered computer system to map network addresses of groups to "group names". . . The address of an entity on a network, including that of a cluster node or a group, is typically obtained in a conventional clustered computer system by accessing a network name server such as a directory name service (DNS) server resident on the network. Thus, should an entity desire to access another entity on a network, the accessing entity typically resolves the name of the entity to be accessed through the network name server, and then sends a message to the network address returned by the server. Thus, in the case of an external access to a group, an entity wishing to send a

request to the group resolves the group name through the network name server, and sends a message to the group address that is returned by the server.

Applicants address the problems associated with this type of access to a clustered computer system at page 3, lines 6-28 of the Application. In brief, access to a group via a group name that has been publicized to an external name server such as a DNS server increases the risk that an unauthorized entity could maliciously access a group. In addition, a DNS server typically can only resolve a group name to a single network address, and as such, it is often not possible to utilize the same publicized group name to access different groups running in different cluster instances executing on the same physical machine (e.g., in the case of a logically partitioned computer).

As Applicants have repeatedly pointed out throughout the prosecution of this Application, claim 1 is directed specifically to the use of a "cluster-private group name" to access a cluster group. The Application specifically defines what makes a group name "cluster-private" at page 7, lines 6-13:

> A group name is cluster-private if the group name cannot be accessed outside of a node that participates in a particular cluster, or cluster "instance". A group name may be cluster-private, for example, if the group name is resident upon a node that participates in a cluster instance, and is accessible by jobs executing on that node, i.e., if the group name is local to the node. A group name may also be cluster-private, for example, if some form of authentication mechanism is utilized to restrict access to the group name only to nodes and/or jobs that can establish participation in a cluster.

Therefore, a key aspect of a "cluster-private group name" is that the name cannot be accessed outside of a node that participates in a cluster.

Chung, on the other hand, is specifically directed to accessing a cluster using a domain name that is "published" to a DNS (col. 7, lines 27-28). As such, even if the domain name were somehow analogized to a group name, it could not be considered to be "cluster-private" in nature. Indeed, the purpose of the *Chung* cluster is to provide client access to a public web site or other cluster service, so it is quite apparent that the domain name published by *Chung* is public, and not private in nature.

It is also important to note that *Chung* does not even disclose the concept of a cluster "group" within the context of Applicants' invention. As described at page 1, line 27 to page 2, line 4 of the Application, a group is defined as a logic entity in a cluster that is used to manage a set of cooperative jobs resident on different nodes of a cluster. In addition, claim 1 specifically recites a "group [that] includes a plurality of members resident respectively on [a] plurality of nodes." There is no disclosure of the use of groups, much less groups identified and accessed via"cluster-private group names," in *Chung*.

As such, *Chung* does not disclose the use of a logical entity known as a group, or the use of a cluster-private group name that cannot be accessed outside of a node that participates in a cluster, as is specifically recited in claim 1 and defined in the specification. Accordingly, claim 1 is novel over *Chung*.

In responding to Applicants' arguments, the Examiner has asserted in the Final Office Action dated November 4, 2005 that Applicants' arguments are not persuasive because Applicants are arguing non-claimed limitations, specifically regarding the concept of a "unique group identifier that cannot be accessed outside of a node that participates in a cluster."

However, Applicants' arguments with regard to a cluster-private group name being an identifier that cannot be accessed outside of a node that participates in a cluster were made for the purpose of explaining the meaning of the claimed term "cluster-private group name." While the Examiner is correct regarding the general proposition that the words of a claim must be given their "plain meaning," MPEP 2111.01 cautions that Applicants are entitled to be their own lexicographers, and that the meaning of a claim term may be interpreted based upon the specification whenever the term is "defined" in the specification. As noted above, the Application specifically states at page 7, lines 6-8 that "[a] group name is cluster-private if the group name cannot be accessed outside of a node that participates in a particular cluster, or cluster 'instance.'" Applicants therefore respectfully submit that the term "cluster-private group name" has in fact been defined in the specification, and as such, the term should be interpreted in a manner that is consistent with this definition. When properly construed, there is no reasonable interpretation of *Chung* that discloses this concept.

Furthermore, irrespective of whether a "cluster-private group name" is required to be interpreted as a "unique group identifier that cannot be accessed outside of a node that participates in a cluster" despite that specific language not being explicitly incorporated into claim 1, it is important to note that claim 1 specifically recites a "cluster-private group name," and the Examiner has never presented any explanation of where the concept of a group name that is "cluster-private" is allegedly disclosed in *Chung*.[1] Put another way, the Examiner has failed to assert any particular "plain meaning" for the concept of a "cluster-private group name," much less a plain meaning that is even arguably disclosed by *Chung*. As noted above, *Chung* discloses a domain name that is specifically published by a DNS server, and is accessible outside of the nodes participating in a cluster. Applicants respectfully submit that a domain name that is published by a DNS server, and is thus accessible by entities that do not participate in a cluster, is not "cluster-private" in nature, and thus cannot form the basis for anticipating the concept of a "cluster-private group name" as is recited in claim 1.

Applicants therefore respectfully submit that *Chung* fails to anticipate claim 1, and the rejection of the claim should be reversed.

Furthermore, given that *Chung* has no appreciation whatsoever for the desirability of privatizing cluster group access, Applicants submit that one of ordinary skill in the art would not be motivated to modify *Chung* to incorporate any such functionality. Applicants therefore respectfully submit that claim 1 is also non-obvious over *Chung*. Reversal of the Examiner's rejection of claim 1, as well as allowance of the claim and of claims 2-14 which depend therefrom, are therefore respectfully requested.

---

[1]Indeed, it should be noted that the Final Office Action is noticeably deficient of critical analysis relating to how *Chung* applies to *any* of the claims at issue. Each rejection follows the basic format of reproducing the text of the rejected claim, followed by a citation to a passage in *Chung* without any discussion of how the passage applies to the claim language. Moreover, the only arguments presented in rebuttal to Applicants' last Response were a terse, stock admonition that limitations not found in the claims are not entitled to patentable weight.

Independent Claims 15, 25 and 26

Similar to claim 1, each of independent claims 15, 25 and 26 recites in part the concept of accessing a group in a clustered computer system using a cluster-private group name. As discussed above in connection with claim 1, this concept is not disclosed or suggested by *Chung*. Accordingly, Applicants respectfully submit that each of these claims is likewise novel and non-obvious over the prior art of record for the same reasons as presented above for claim 1. Reconsideration and allowance of independent claims 15, 25 and 26, as well as of claims 16-24 and 27 which depend respectively therefrom, are therefore respectfully requested.

Dependent Claims 2 and 16

Claim 2 (to which claim 16 is similar) depends from claim 1 and additionally recites generating the access request with a user job resident on the first node. In rejecting the claim, the Examiner once again cites col. 7, lines 38-52 of *Chung*, but without applying the passage to the specific language of the claim. Applicants can find no applicability of this passage to the specific language recited in the claim, and as such, Applicants submit the Examiner has failed to establish anticipation of this claim by *Chung*. The passage in *Chung*, for example, discloses a request being issued by a client that accesses a cluster, while claim 2 refers to a user job resident on a first node which itself is part of a cluster. Accordingly, Applicants submit that the Examiner has not met the burden of establishing anticipation of claim 2 by *Chung*. Reversal of the Examiner's rejection of claim 2 (and of claim 16 that is similar thereto) is therefore respectfully requested.

Dependent Claims 3 and 17

Claim 3 (to which claim 17 is similar) depends from claim 2 and additionally recites forwarding the access request to a clustering infrastructure resident in the first node via a call from the user job. In rejecting the claim, the Examiner once again cites col. 7, lines 38-52 of *Chung*, but without applying the passage to the specific language of the claim. Applicants can find no applicability of this passage to the specific language recited in the claim, and as such, Applicants submit the Examiner has failed to establish anticipation of this claim by *Chung*. The passage in *Chung*, for example, is completely silent with respect to the concept of a clustering

infrastructure resident on a node of a cluster. Accordingly, Applicants submit that the Examiner has not met the burden of establishing anticipation of claim 3 by *Chung*. Reversal of the Examiner's rejection of claim 3 (and of claim 17 that is similar thereto) is therefore respectfully requested.

Dependent Claims 4 and 18

Claim 4 (to which claim 18 is similar) depends from claim 1 and additionally recites generating the access request with a user job resident on a second node in the plurality of nodes, and processing the access request with a proxy job resident on the second node by communicating the access request to the first node. In rejecting the claim, the Examiner cites col. 7, line 62 to col. 8, line 15 of *Chung*, but without applying the passage to the specific language of the claim. Applicants can find no applicability of this passage to the specific language recited in the claim, and as such, Applicants submit the Examiner has failed to establish anticipation of this claim by *Chung*. The passage in *Chung*, for example, refers to a router that routes packets to nodes in a cluster. However, there is no discussion of a user job or a proxy job anywhere in the passage. Accordingly, Applicants submit that the Examiner has not met the burden of establishing anticipation of claim 4 by *Chung*. Reversal of the Examiner's rejection of claim 4 (and of claim 18 that is similar thereto) is therefore respectfully requested.

Dependent Claim 5

Claim 5 depends from claim 4 and additionally recites that the proxy job is a member of a cluster control group. Claim 5 also recites forwarding the access request from the user job to the proxy job, and forwarding the access request from the proxy job to a clustering infrastructure resident in the second node via a call from the proxy job. In rejecting the claim, the Examiner again cites col. 7, line 62 to col. 8, line 15 of *Chung*, but without applying the passage to the specific language of the claim. Applicants can find no applicability of this passage to the specific language recited in the claim, and as such, Applicants submit the Examiner has failed to establish anticipation of this claim by *Chung*. The passage in *Chung*, for example, refers to a router that routes packets to nodes in a cluster. However, there is no discussion of a user job, a proxy job, or

a clustering infrastructure anywhere in the passage. Accordingly, Applicants submit that the Examiner has not met the burden of establishing anticipation of claim 5 by *Chung*. Reversal of the Examiner's rejection of claim 5 is therefore respectfully requested.

Dependent Claims 6 and 19

Claim 6 (to which claim 19 is similar) depends from claim 1 and additionally recites retrieving the cluster-private group name with a user job by accessing a cluster-private data structure. In rejecting the claim, the Examiner cites col. 8, lines 16-48 of *Chung*, but without applying the passage to the specific language of the claim. Applicants can find no applicability of this passage to the specific language recited in the claim, and as such, Applicants submit the Examiner has failed to establish anticipation of this claim by *Chung*. The passage in *Chung*, for example, is completely silent with respect to any form of data structure, much less a "cluster-private" data structure. Accordingly, Applicants submit that the Examiner has not met the burden of establishing anticipation of claim 6 by *Chung*. Reversal of the Examiner's rejection of claim 6 (and of claim 19 that is similar thereto) is therefore respectfully requested.

Dependent Claims 7 and 20

Claim 7 (to which claim 20 is similar) depends from claim 6 and additionally recites that the cluster-private data structure is resident on the same node as the user job. In rejecting the claim, the Examiner again cites col. 8, lines 16-48 of *Chung*, but without applying the passage to the specific language of the claim. Applicants can find no applicability of this passage to the specific language recited in the claim, and as such, Applicants submit the Examiner has failed to establish anticipation of this claim by *Chung*. The passage in *Chung*, for example, is completely silent with respect to any form of data structure, much less a "cluster-private" data structure. Furthermore, the passage does not disclose any relationship between a data structure and a user job as being resident on the same node. Accordingly, Applicants submit that the Examiner has not met the burden of establishing anticipation of claim 7 by *Chung*. Reversal of the Examiner's rejection of claim 7 (and of claim 20 that is similar thereto) is therefore respectfully requested.

Dependent Claims 8 and 21

Claim 8 (to which claim 21 is similar) depends from claim 7 and additionally recites that the cluster-private data structure is accessible only from the node upon which the cluster-private data structure is resident. In rejecting the claim, the Examiner again cites col. 8, lines 16-48 of *Chung*, but without applying the passage to the specific language of the claim. Applicants can find no applicability of this passage to the specific language recited in the claim, and as such, Applicants submit the Examiner has failed to establish anticipation of this claim by *Chung*. The passage in *Chung*, for example, is completely silent with respect to any form of data structure, much less a "cluster-private" data structure. Furthermore, the passage does not disclose constraint on the accessibility of a data structure. Accordingly, Applicants submit that the Examiner has not met the burden of establishing anticipation of claim 8 by *Chung*. Reversal of the Examiner's rejection of claim 8 (and of claim 21 that is similar thereto) is therefore respectfully requested.

Dependent Claim 9

Claim 9 depends from claim 8 and additionally recites that the cluster-private data structure is accessible only by jobs that are resident on the node upon which the cluster-private data structure is resident. In rejecting the claim, the Examiner cites col. 8, line 50 to col. 9, line 23 of *Chung*, but without applying the passage to the specific language of the claim. Applicants can find no applicability of this passage to the specific language recited in the claim, and as such, Applicants submit the Examiner has failed to establish anticipation of this claim by *Chung*. The passage in *Chung*, for example, is completely silent with respect to any form of data structure, much less a "cluster-private" data structure. Furthermore, the passage does not disclose constraint on the accessibility of a data structure based upon the location of a job. Accordingly, Applicants submit that the Examiner has not met the burden of establishing anticipation of claim 9 by *Chung*. Reversal of the Examiner's rejection of claim 9 is therefore respectfully requested.

Dependent Claim 10

Claim 10 is not separately argued.

Dependent Claims 11 and 22

Claim 11 (to which claim 22 is similar) depends from claim 10 and additionally recites that initiating the group operation further comprises accessing a group address data structure to retrieve a plurality of network addresses associated with the cluster-private group name, and that distributing messages to the plurality of group members includes sending a message to each of the plurality of network addresses. In rejecting the claim, the Examiner cites col. 9, lines 23-50 of *Chung*, but without applying the passage to the specific language of the claim. Applicants can find no applicability of this passage to the specific language recited in the claim, and as such, Applicants submit the Examiner has failed to establish anticipation of this claim by *Chung*. The passage in *Chung*, for example, is completely silent with respect to the concept of a group address data structure, or of distributing messages to group members using the group address data structure. Accordingly, Applicants submit that the Examiner has not met the burden of establishing anticipation of claim 11 by *Chung*. Reversal of the Examiner's rejection of claim 11 (and of claim 22 that is similar thereto) is therefore respectfully requested.

Dependent Claim 12

Claim 12 depends from claim 1 and additionally recites that initiating the group operation is performed by a clustering infrastructure resident on the first node. In rejecting the claim, the Examiner cites col. 9, lines 23-50 of *Chung*, but without applying the passage to the specific language of the claim. Applicants can find no applicability of this passage to the specific language recited in the claim, and as such, Applicants submit the Examiner has failed to establish anticipation of this claim by *Chung*. The passage in *Chung*, for example, is completely silent with respect to any form of clustering infrastructure, much less one that initiates a group operation. Accordingly, Applicants submit that the Examiner has not met the burden of establishing anticipation of claim 12 by *Chung*. Reversal of the Examiner's rejection of claim 12 is therefore respectfully requested.

Dependent Claims 13 and 23

Claim 13 (to which claim 23 is similar) depends from claim 12 and additionally recites that initiating the group operation includes retrieving with the clustering infrastructure a plurality of addresses that are mapped to the cluster-private group name in a data structure that is local to the clustering infrastructure. In rejecting the claim, the Examiner again cites col. 9, lines 23-50 of *Chung*, but without applying the passage to the specific language of the claim. Applicants can find no applicability of this passage to the specific language recited in the claim, and as such, Applicants submit the Examiner has failed to establish anticipation of this claim by *Chung*. The passage in *Chung*, for example, is completely silent with respect to the concept of a group name that is local to a clustering infrastructure. Accordingly, Applicants submit that the Examiner has not met the burden of establishing anticipation of claim 13 by *Chung*. Reversal of the Examiner's rejection of claim 13 (and of claim 23 that is similar thereto) is therefore respectfully requested.

Dependent Claims 14 and 24

Claim 14 (to which claim 24 is similar) depends from claim 1 and additionally recites that initiating the group operation includes locally resolving on the first node a mapping between the cluster-private group name and a plurality of addresses associated with at least the subset of the plurality of nodes. In rejecting the claim, the Examiner cites col. 7, line 15 to col. 8, line 51 of *Chung*, but without applying the passage to the specific language of the claim. Applicants can find no applicability of this passage to the specific language recited in the claim, and as such, Applicants submit the Examiner has failed to establish anticipation of this claim by *Chung*. The passage in *Chung*, for example, is completely silent with respect to the concept of "locally resolving" a mapping between a cluster-private group name and addresses associated with nodes. Accordingly, Applicants submit that the Examiner has not met the burden of establishing anticipation of claim 14 by *Chung*. Reversal of the Examiner's rejection of claim 14 (and of claim 24 that is similar thereto) is therefore respectfully requested.

Dependent Claim 27

Claim 27 is not argued separately.

CONCLUSION

In conclusion, Applicants respectfully request that the Board reverse the Examiner's rejections of claims 1-27, and that the Application be passed to issue. If there are any questions regarding the foregoing, please contact the undersigned at 513/241-2324. Moreover, if any other charges or credits are necessary to complete this communication, please apply them to Deposit Account 23-3000.

Respectfully submitted,

WOOD, HERRON & EVANS, L.L.P.

Date: June 21, 2006

By: /Scott A. Stinebruner/
Scott A. Stinebruner
Reg. No. 38,323

2700 Carew Tower
441 Vine Street
Cincinnati, Ohio 45202
(513) 241-2324 - Telephone
(513) 241-6234 - Facsimile

VIII. CLAIMS APPENDIX: CLAIMS ON APPEAL (S/N 09/845,596)

1. (Original) A method of accessing a group in a clustered computer system, wherein the clustered computer system includes a plurality of nodes, and wherein the group includes a plurality of members resident respectively on the plurality of nodes, the method comprising:

(a) receiving an access request on a first node in the plurality of nodes, wherein the access request identifies a cluster-private group name associated with the group; and

(b) processing the access request on the first node to initiate a group operation on at least a subset of the plurality of nodes that map to the cluster-private group name.

2. (Original) The method of claim 1, further comprising generating the access request with a user job resident on the first node.

3. (Original) The method of claim 2, further comprising forwarding the access request to a clustering infrastructure resident in the first node via a call from the user job.

4. (Original) The method of claim 1, further comprising:

(a) generating the access request with a user job resident on a second node in the plurality of nodes; and

(b) processing the access request with a proxy job resident on the second node by communicating the access request to the first node.

5. (Original) The method of claim 4, wherein the proxy job is a member of a cluster control group, the method further comprising:

(a) forwarding the access request from the user job to the proxy job; and

(b) forwarding the access request from the proxy job to a clustering infrastructure resident in the second node via a call from the proxy job.

6. (Original) The method of claim 1, further comprising retrieving the cluster-private group name with a user job by accessing a cluster-private data structure.

7. (Original) The method of claim 6, wherein the cluster-private data structure is resident on the same node as the user job.

8. (Original) The method of claim 7, wherein the cluster-private data structure is accessible only from the node upon which the cluster-private data structure is resident.

9. (Original) The method of claim 8, wherein the cluster-private data structure is accessible only by jobs that are resident on the node upon which the cluster-private data structure is resident.

10. (Original) The method of claim 1, wherein initiating the group operation comprises distributing messages to a plurality of group members resident on the nodes that map to the cluster-private group name.

11. (Original) The method of claim 10, wherein initiating the group operation further comprises accessing a group address data structure to retrieve a plurality of network addresses associated with the cluster-private group name, wherein distributing messages to the plurality of group members includes sending a message to each of the plurality of network addresses.

12. (Original) The method of claim 1, wherein initiating the group operation is performed by a clustering infrastructure resident on the first node.

13. (Original) The method of claim 12, wherein initiating the group operation includes retrieving with the clustering infrastructure a plurality of addresses that are mapped to the cluster-private group name in a data structure that is local to the clustering infrastructure.

14. (Original) The method of claim 1, wherein initiating the group operation includes locally resolving on the first node a mapping between the cluster-private group name and a plurality of addresses associated with at least the subset of the plurality of nodes.

15. (Original) An apparatus, comprising:

(a) a memory accessible by a first node among a plurality of nodes in a clustered computer system; and

(b) a program resident in the memory and executed by the first node, the program configured to access a group that includes a plurality of members resident respectively on the plurality of nodes by receiving an access request that identifies a cluster-private group name associated with the group, and processing the access request to initiate a group operation on at least a subset of the plurality of nodes that map to the cluster-private group name.

16. (Original) The apparatus of claim 15, further comprising a user job configured to generate the access request.

17. (Original) The apparatus of claim 16, wherein the program comprises a clustering infrastructure resident on the first node.

18. (Original) The apparatus of claim 17, further comprising a proxy job configured to forward the access request from the user job to the clustering infrastructure.

19. (Original) The apparatus of claim 15, further comprising:

(a) a cluster-private data structure configured to store the cluster-private group name; and

(b) a user job configured to access the cluster-private data structure to retrieve the cluster-private group name and generate the access request therefrom.

20. (Original) The apparatus of claim 19, wherein the cluster-private data structure is resident on the same node as the user job.

21. (Original) The apparatus of claim 20, wherein the cluster-private data structure is accessible only from the node upon which the cluster-private data structure is resident.

22. (Original) The apparatus of claim 15, further comprising a group address data structure configured to store a plurality of network addresses associated with the cluster-private group name, wherein the program is configured to initiate the group operation by accessing the group address data structure to retrieve the plurality of network addresses and sending a message to each of the plurality of network addresses.

23. (Original) The apparatus of claim 22, wherein the program comprises a clustering infrastructure, and wherein the group address data structure is local to the clustering infrastructure.

24. (Once Amended) The apparatus of claim 15, wherein the program is further configured to process the access request by locally resolving on the first node a mapping between the cluster-private group name and a plurality of addresses associated with at least the subset of the plurality of nodes.

25. (Original) A clustered computer system, comprising:

(a) a plurality of nodes coupled to one another over a network;

(b) a group including a plurality of members resident respectively on the plurality of nodes; and

(c) a program resident in a first node among the plurality of nodes and configured to access the group by receiving an access request that identifies a cluster-private group name associated with the group, and processing the access request to initiate a group operation on at least a subset of the plurality of nodes that map to the cluster-private group name.

26. (Original) A program product, comprising:

(a) a program resident in the memory and executed by a first node among a plurality of nodes in a clustered computer system, the program configured to access a group that includes a plurality of members resident respectively on the plurality of nodes by receiving an access request that identifies a cluster-private group name associated with

the group, and processing the access request to initiate a group operation on at least a subset of the plurality of nodes that map to the cluster-private group name; and

(b) a signal bearing medium bearing the program.

27. (Original) The program product of claim 26, wherein the signal bearing medium includes at least one of a transmission medium and a recordable medium

IX. EVIDENCE APPENDIX

09/845,596

None.

X. RELATED PROCEEDINGS APPENDIX

09/845,596

None.